



08/754429

- 1 -

~~GENERAL PURPOSE, PROGRAMMABLE MEDIA PROCESSOR~~

AI >

5 A Microfiche Appendix consisting of 4 sheets (387 total frames) of  
microfiche is included in this application. The Microfiche Appendix contains  
material which is subject to copyright protection. The copyright owner has no  
objection to the facsimile reproduction by any one of the Microfiche Appendix, as  
it appears in the Patent and Trademark Office patent files or records, but  
10 otherwise reserves all copyright rights whatsoever.

Field of the Invention

This invention relates to the field of communications processing,  
and more particularly, to a method and apparatus for real-time processing of  
15 multi-media digital communications.

Background of the Invention

Optical fiber and discs have made the transmission and storage of  
digital information both cheaper and easier than older analog technologies. An  
20 improved system for digital processing of media data streams is necessary in order  
to realize the full potential of these advanced media.

For the past century, telephone service delivered over copper  
twisted pair has been the lingua franca of communications. Over the next century,  
broadband services delivered over optical fiber and coax will more completely

1 -  
Z

fulfill the human need for sensory information by supplying voice, video, and data at rates of about 1,000 times greater than narrow band telephony. Current general-purpose microprocessors and digital signal processors ("DSPs") can handle digital voice, data, and images at narrow band rates, but they are way too slow for processing media data at broadband rates.

This shortfall in digital processing of broadband media is currently being addressed through the design of many different kinds of application-specific integrated circuits ("ASICs"). For example, a prototypical broadband device such as a cable modem modulates and demodulates digital data at rates up to 45 Mbits/sec within a single 6 MHz cable channel (as compared to rates of 28.8 Kbits/sec within a 6 KHz channel for telephone modems) and transcodes it onto a 10/100baseT connection to a personal computer ("PC") or workstation. Current cable modems thus receive data from a coaxial cable connection through a chain of specialized ASIC devices in order to accomplish Quadrature Amplitude Modification ("QAM") demodulation, Reed-Solomon error correction, packet filtering, Data Encryption Standard ("DES") decryption, and Ethernet protocol handling. The cable modems also transmit data to the coaxial cable link through a second chain of devices to achieve DES encryption, Reed-Solomon block encoding, and Quaternary Phase Shift Keying ("QPSK") modulation. In these environments, a general-purpose processor is usually required as well in order to perform initialization, statistics collection, diagnostics, and network management functions.

The ASIC approach to media processing has three fundamental flaws: cost, complexity, and rigidity. The combined silicon area of all the specialized ASIC devices required in the cable modem, for example, results in a component cost incompatible with the per subscriber price target for a cable service. The cable plant itself is a very hostile service environment, with noise ingress, reflections, nonlinear amplifiers, and other channel impairments, especially when viewed in the upstream direction. Telephony modems have developed an elaborate hierarchy of algorithms implemented in DSP software, with automatic reduction of data rates from 28.8Kbits/sec to 19.6Kbits/sec,

3

14.4Kbits/sec, or much lower rates as needed to accommodate noise, echoes, and other impairments in the copper plant. To implement similar algorithms on an ASIC-based broadband modem is far more complex to achieve in software.

5        These problems of cost, complexity, and rigidity are compounded further in more complete broadband devices such as digital set-top boxes, multimedia PCs, or video conferencing equipment, all of which go beyond the basic radio frequency ("RF") modem functions to include a broad range of audio and video compression and decoding algorithms, along with remote control and graphical user interfaces. Software for these devices must control what amounts to  
10       a heterogeneous multi-processor, where each specialized processor has a different, and usually eccentric or primitive, programming environment. Even if these programming environments are mastered, the degree of programmability is limited. For example, Motion Picture Expert Group-I ("MPEG-I") chips manufactured by AT&T Corporation will not implement advances such as fractal-  
15       and wavelet-based compression algorithms, but these chips are not readily software upgradeable to the MPEG-II standard. A broadband network operator who leases an MPEG ASIC-based product is therefore at risk of having to continuously upgrade his system by purchasing significant amounts of new hardware just to track the evolution of MPEG standards.

20       The high cost of ASIC-based media processing results from inefficiencies in both memory and logic. A typical ASIC consists of a multiplicity of specialized logic blocks, each with a small memory dedicated to holding the data which comprises the working set for that block. The silicon area of these multiple small memories is further increased by the overhead of multiple decoders,  
25       sense amplifiers, write drivers, etc. required for each logic block. The logic blocks are also constrained to operate at frequencies determined by the internal symbol rates of broadband algorithms in order to avoid additional buffer memories. These frequencies typically differ from the optimum speed-area operating point of a given semiconductor technology. Interconnect and  
30       synchronization of the many logic and memory blocks are also major sources of overhead in the ASIC approach.

4 - 3 -

The disadvantages of the prior ASIC approach can be over come by a single unified media processor. The cost advantages of such a unified processor can be achieved by gathering all the many ASIC functions of a broadband media product into a single integrated circuit. Cost reduction is further increased by  
5 reducing the total memory area of such a circuit by replacing the multiplicity of small ASIC memories with a single memory hierarchy large enough to accommodate the sum total of all the working sets, and wide enough to supply the aggregate bandwidth needs of all the logic blocks. Additionally, the logic block interconnect circuitry to this memory hierarchy may be streamlined by providing a  
10 generally programmable switching fabric. Many of the logic blocks themselves can also be replaced with a single multi-precision arithmetic unit, which can be internally partitioned under software control to perform addition, multiplication, division, and other integer and floating point arithmetic operations on symbol streams of varying widths, while sustaining the full data throughput of the memory  
15 hierarchy. The residue of logic blocks that perform operations that are neither arithmetic or permutation group oriented can be replaced with an extended math unit that supports additional arithmetic operations such as finite field, ring, and table lookup, while also sustaining the full data throughput of the memory hierarchy.

20 The above multi-precision arithmetic, permutation switch, and extended math operations can then be organized as machine instructions that transfer their operands to and from a single wide multi-ported register file. These instructions can be further supplemented with load/store instructions that transfer register data to and from a data buffer/cache static random access memory  
25 ("SRAM") and main memory dynamic random access memories ("DRAMs"), and with branch instructions that control the flow of instructions executed from an instruction buffer/cache SRAM. Extensions to the load/store instructions can be made for synchronization, and to branch instructions for protected gateways, so that multiple threads of execution for audio, video, radio, encryption, networking,  
30 etc. can efficiently and securely share memory and logic resources of a unified machine operating near the optimum speed-area point of the target semiconductor

process. The data path for such a unified media processor can interface to a high speed input/output ("I/O") subsystem that moves media streams across ultra-high bandwidth interfaces to external storage and I/O.

5 Such a device would incorporate all of the processing capabilities of the specialized multi-ASIC combination into a single, unified processing device. The unified processor would be agile and capable of reprogramming through the transmission of new programs over the communication medium. This programmable, general purpose device is thus less costly than the specialized processor combination, easier to operate and reprogram and can be installed or  
10 applied in many differing devices and situations. The device may also be scalable to communications applications that support vast numbers of users through massively parallel distributed computing.

It is therefore an object of this invention to process media data streams by executing operations at very high bandwidth rates.

15 It is also an object of this invention to unify the audio, video, radio, graphics, encryption, authentication, and networking protocols into a single instruction stream.

It is also an object of this invention to achieve high bandwidth rates in a unified processor that is easy to program and more flexible than a  
20 heterogeneous combination of special purpose processors.

It is a further object of the invention to support high level mathematical processing in a unified media processor, including finite group, finite field, finite ring and table look-up operations, all at high bandwidth rates.

It is yet a further object of the invention to provide a unified media  
25 processor that can be replicated into a multi-processor system to support a vast array of users.

It is yet another object of this invention to allow for massively parallel systems within the switching fabric to support very large numbers of subscribers and services.

30 It is also an object of the invention to provide a general purpose programmable processor that could be employed at all points in a network.

It is a further object of this invention to sustain very high bandwidth rates to arbitrarily large memory and input/output systems.

#### Summary of the Invention

5                   In view of the above, there is provided a system for media processing that maintains substantially peak data throughput in the execution and transmission of multiple media data streams. The system includes in one aspect a general purpose, programmable media processor, and in another aspect includes a method for receiving, processing and transmitting media data streams. The  
10                  general purpose, programmable media processor of the invention further includes an execution unit, high bandwidth external interface, and can be employed in a parallel multi-processor system.

                  According to the apparatus of the invention, an execution unit is provided that maintains substantially peak data throughput in the unified execution  
15                  of multiple media data streams. The execution unit includes a data path, and a multi-precision arithmetic unit coupled to the data path and capable of dynamic partitioning based on the elemental width of data received from the data path. The execution unit also includes a switch coupled to the data path that is programmable to manipulate data received from the data path and provide data streams to the  
20                  data path. An extended mathematical element is also provided, which is coupled to the data path and programmable to implement additional mathematical operations at substantially peak data throughput. In a preferred embodiment of the execution unit, at least one register file is coupled to the data path.

                  According to another aspect of the invention, a general purpose  
25                  programmable media processor is provided having an instruction path and a data path to digitally process a plurality of media data streams. The media processor includes a high bandwidth external interface operable to receive a plurality of data of various sizes from an external source and communicate the received data over the data path at a rate that maintains substantially peak operation of the media  
30                  processor. At least one register file is included, which is configurable to receive and store data from the data path and to communicate the stored data to the data

path. A multi-precision execution unit is coupled to the data path and is dynamically configurable to partition data received from the data path to account for the elemental symbol size of the plurality of media streams, and is programmable to operate on the data to generate a unified symbol output to the data path.

According to the preferred embodiment of the media processor, means are included for moving data between registers and memory by performing load and store operations, and for coordinating the sharing of data among a plurality of tasks by performing synchronization operations based upon instructions and data received by the execution unit. Means are also provided for securely controlling the sequence of execution by performing branch and gateway operations based upon instructions and data received by the execution unit. A memory management unit operable to retrieve data and instructions for timely and secure communication over the data path and instruction path respectively is also preferably included in the media processor. The preferred embodiment also includes a combined instruction cache and buffer that is dynamically allocated between cache space and buffer space to ensure real-time execution of multiple media instruction streams, and a combined data cache and buffer that is dynamically allocated between cache space and buffer space to ensure real-time response for multiple media data streams.

In another aspect of the invention, a high bandwidth processor interface for receiving and transmitting a media stream is provided having a data path operable to transmit media information at sustained peak rates. The high bandwidth processor interface includes a plurality of memory controllers coupled in series to communicate stored media information to and from the data path, and a plurality of memory elements coupled in parallel to each of the plurality of memory controllers for storing and retrieving the media information. In the preferred embodiment of the high bandwidth processor interface, the plurality of memory controllers each comprise a paired link disposed between each memory controller, where the paired links each transmit and receive plural bits of data and have differential data inputs and outputs and a differential clock signal.

8

Yet another aspect of the invention includes a system for unified media processing having a plurality of general purpose media processors, where each media processor is operable at substantially peak data rates and has a dynamically partitioned execution unit and a high bandwidth interface for communicating to memory and input/output elements to supply data to the media processor at substantially peak rates. A bi-directional communication fabric is provided, to which the plurality of media processors are coupled, to transmit and receive at least one media stream comprising presentation, transmission, and storage media information. The bi-directional communication fabric preferably comprises a fiber optic network, and a subset of the plurality of media processors comprise network servers.

According to yet another aspect of the invention, a parallel multi-media processor system is provided having a data path and a high bandwidth external interface coupled to the data path and operable to receive a plurality of data of various sizes from an external source and communicate the received data at a rate that maintains substantially peak operation of the parallel multi-processor system. A plurality of register files, each having at least one register coupled to the data path and operable to store data, are also included. At least one multi-precision execution unit is coupled to the data path and is dynamically configurable to partition data received from the data path to account for the elemental symbol size of the plurality of media streams, and is programmable to operate in parallel on data stored in the plurality of register files to generate a unified symbol output for each register file.

According to the method of the invention, unified streams of media data are processed by receiving a stream of unified media data including presentation, transmission and storage information. The unified stream of media data is dynamically partitioned into component fields of at least one bit based on the elemental symbol size of data received. The unified stream of media data is then processed at substantially peak operation.

In one aspect of the invention, the unified stream of media data is processed by storing the stream of unified media data in a general register file.



Multi-precision arithmetic operations can then be performed on the stored stream of unified media data based on programmed instructions, where the multi-precision arithmetic operations include Boolean, integer and floating point mathematical operations. The component fields of unified media data can then be manipulated  
5 based on programmed instructions that implement copying, shifting and re-sizing operations. Multi-precision mathematical operations can also be performed on the stored stream of unified media data based on programmed instructions, where the mathematical operations including finite group, finite field, finite ring and table look-up operations. Instruction and data pre-fetching are included to fill  
10 instruction and data pipelines, and memory management operations can be performed to retrieve instructions and data from external memory. The instructions and data are preferably stored in instruction and data cache/buffers, in which buffer storage in the instruction and data cache/buffers is dynamically allocated to ensure real-time execution.

15 Other aspects of the invention include a method for achieving high bandwidth communications between a general purpose media processor and external devices by providing a high bandwidth interface disposed between the media processor and the external devices, in which the high bandwidth interface comprises at least one uni-directional channel pair having an input port and an  
20 output port. A plurality of media data streams, comprising component fields of various sizes, are transmitted and received between the media processor and the external devices at a rate that sustains substantially peak data throughput at the media processor. A method for processing streams of media data is also included that provides a bi-directional communications fabric for transmitting and receiving  
25 at least one stream of media data, where the at least one stream of media data comprises presentation, transmission and storage information. At least one programmable media processor is provided within the communications network for receiving, processing and transmitting the at least one stream of unified media data over the bi-directional communications fabric.

30 The general purpose, programmable media processor of the invention combines in a single device all of the necessary hardware included in the

10

specialized processor combinations to process and communicate digital media data streams in real-time. The general purpose, programmable media processor is therefore cheaper and more flexible than the prior approach to media processing. The general purpose, programmable media processor is thus more susceptible to  
5 incorporation within a massively parallel processing network of general purpose media processors that enhance the ability to provide real-time multi-media communications to the masses.

These features are accomplished by deploying server media processors and client media processors throughout the network. Such a network  
10 provides a seamless, global media super-computer which allows programmers and network owners to virtualize resources. Rather than restrictively accessing only the memory space and processing time of a local resource, the system allows access to resources throughout the network. In small access points such as wireless devices, where very little memory and processing logic is available due to  
15 limited battery life, the system is able to draw upon the resources of a homogeneous multi-computer system.

The invention also allows network owners the facility to track standards and to deploy new services by broadcasting software across the network rather than by instituting costly hardware upgrades across the whole network.  
20 Broadcasting software across the network can be performed at the end of an advertisement or other program that is broadcasted nationally. Thus, services can be advertised and then transmitted to new subscribers at the end of the advertisement.

These and other features and advantages of the invention will be  
25 apparent upon consideration of the following detailed description of the presently preferred embodiments of the invention, taken in conjunction with the appended drawings.

#### **Brief Description of the Drawings**

30 FIG. 1 is a block diagram of a broad band media computer employing the general purpose, programmable media processor of the invention;

FIG. 2 is a block diagram of a global media processor employing multiple general purpose media processors according to the invention;

FIG. 3 is an illustration of the digital bandwidth spectrum for telecommunications, media and computing communications;

5                   FIG. 4 is the digital bandwidth spectrum shown in FIG. 3 taking into account the bandwidth overhead associated with compressed video techniques;

FIG. 5 is a block diagram of the current specialized processor solution for mass media communication, where FIG. 5(a) shows the current distributed system, and FIG. 5(b) shows a possible integrated approach;

10                   FIG. 6 is a block diagram of two presently preferred general purpose media processors, where FIG. 6(a) shows a distributed system and FIG. 6(b) shows an integrated media processor;

FIG. 7 is a block diagram of the presently preferred structure of a general purpose, programmable media processor according to the invention;

15                   FIG. 8 is a drawing consisting of visual illustrations of the various group operations provided on the media processor, where FIG. 8(a) illustrates the group expand operation, FIG. 8(b) illustrates the group compress or extract operation, FIG. 8(c) illustrates the group deal and shuffle operations, FIG. 8(d) illustrates the group swizzle operation and FIG. 8(e) illustrates the various group  
20                   permute operations;

FIG. 9 shows the preferred instruction and data sizes for the general purpose, programmable media processor, where FIG. 9(a) is an illustration of the various instruction formats available on the general purpose, programmable media processor, FIG. 9(b) illustrates the various floating-point data sizes available on  
25                   the general purpose media processor, and FIG. 9(c) illustrates the various fixed-point data sizes available on the general purpose media processor;

FIG. 10 is an illustration of a presently preferred memory management unit included in the general purpose processor shown in FIG. 7, where FIG. 10(a) is a translation block diagram and FIG. 10(b) illustrates the  
30                   functional blocks of the transaction lookaside buffer;

FIG. 11 is an illustration of a super-string pipeline technique;

11

FIG. 12 is an illustration of the presently preferred super-spring pipeline technique;

FIG. 13 is a block diagram of a single memory channel for communication to the general purpose media processor shown in FIG. 7;

5           FIG. 14 is an illustration of the presently preferred connection of standard memory devices to the preferred memory interface;

FIG. 15 is a block diagram of the input/output controller for use with the memory channel shown in FIG. 13;

10           FIG. 16 is a block diagram showing multiple memory channels connected to the general purpose media processor shown in FIG. 7, where FIG. 16(a) shows a two-channel implementation and FIG. 16(b) illustrates a twelve-channel embodiment;

FIG. 17 illustrates the presently preferred packet communications protocol for use over the memory channel shown in FIG. 13;

15           FIG. 18 shows a multi-processor configuration employing the general purpose media processor shown in FIG. 7, where FIG. 18(a) shows a linear processor configuration, FIG. 18(b) shows a processor ring configuration, and FIG. 18(c) shows a two-dimensional processor configuration; and

20           FIG. 19 shows a presently preferred multi-chip implementation of the general purpose, programmable media processor of the invention.

#### **Detailed Description of the Presently Preferred Embodiments**

Referring to the drawings, where like-reference numerals refer to like elements throughout, a broad band microcomputer 10 is provided in FIG. 1.

25           The broad band microcomputer 10 consists essentially of a general purpose media processor 12. As will be described in more detail below, the general purpose media processor 12 receives, processes and transmits media data streams in a bi-directional manner from upstream network components to downstream devices. In general, media data streams received from upstream network components can

30           comprise any combination of audio, video, radio, graphics, encryption, authentication, and networking information. As those skilled in the art will

appreciate, however, the general purpose media processor 12 is in no way limited to receiving, processing and transmitting only these types of media information. The general purpose media processor 12 of the invention is capable of processing any form of digital media information without departing from the spirit and  
5 essential scope of the invention.

#### System Configuration

In the preferred embodiment of the invention shown in FIG. 1, media data streams are communicated to the media processor 12 from several  
10 sources. Ideally, unified media data streams are received and transmitted by the general purpose media processor 12 over a fiber optic cable network 14. As will be described in more detail below, although a fiber optic cable network is preferred, the presently existing communications network in the United States consists of a combination of fiber optic cable, coaxial cable and other transmission  
15 media. Consequently, the general purpose media processor 12 can also receive and transmit media data streams over coaxial cable 14 and traditional twisted pair wire connections 16. The specific communications protocol employed over the twisted pair 16, whether POTS, ISDN or ADSL, is not essential; all protocols are supported by the broad band microcomputer 10. The details of these protocols are  
20 generally known to those skilled in the art and no further discussion is therefore needed or provided herein.

Another form of upstream network communication is through a satellite link 18. The satellite link 18 is typically connected to a satellite receiver  
20. The satellite receiver 20 comprises an antenna, usually in the form of a  
25 satellite dish, and amplification circuitry. The details of such satellite communications are also generally known in the art, and further detail is therefore not provided or included herein.

As described above, the general purpose media processor 12 communicates in a bi-directional manner to receive, process and transmit media  
30 data streams to and from downstream devices. As shown in FIG. 1, downstream communication preferably takes place in at least two forms. First, media data

streams can be communicated over a bi-directional local network 22. Various types of local networks 22 are generally known in the art and many different forms exist. The general purpose media processor 12 is capable of communicating over any of these local networks 22 and the particular type of network selected is implementation specific.

The local network 22 is preferably employed to communicate between the unified processor 12, and audio/visual devices 24 or other digital devices 26. Presently preferred examples of audio/visual devices 24 include digital cable television, video-on-demand devices, electronic yellow pages services, integrated message systems, video telephones, video games and electronic program guides. As those skilled in the art will appreciate, other forms of audio/video devices are contemplated within the spirit and scope of the invention. Presently preferred embodiments of other digital devices 26 for communication with the general purpose media processor 12 include personal computers, television sets, work stations, digital video camera recorders, and compact disc read-only memories. As those skilled in the art will also appreciate, further digital devices 26 are contemplated for communication to the general purpose media processor 12 without departing from the spirit and scope of the invention.

Second, the general purpose media processor preferably also communicates with downstream devices over a wireless network 28. In the presently preferred embodiment of the invention, wireless devices for communication over the wireless network 28 can comprise either remote communication devices 30 or remote computing devices 32. Presently preferred embodiments of the remote communications devices 30 include cordless telephones and personal communicators. Presently preferred embodiments of the remote computing devices 32 include remote controls and telecommunicating devices. As those skilled in the art will appreciate, other forms of remote communication devices 30 and remote computing devices 32 are capable of communication with the general purpose media processor 12 without departing from the spirit and scope of the invention. An agile digital radio (not shown) that incorporates a

15  
14 -

general purpose media processor 12 may be used to communicate with these wireless devices.

#### Network Configuration

5 Referring now to FIG. 2, the general purpose media processor 12 is preferably disposed throughout a digital communications network 38. In order to enable communication among large and small businesses, residential customers and mobile users, the network 38 can consist of a combination of many individual sub-networks comprised of three main forms of interconnection. The trunk and main  
10 branches of the network 38 preferably employ fiber optic cable 40 as the preferred means of interconnection. Fiber optic cable 40 is used to connect between general purpose media processors 12 disposed as network servers 46 or large business installations 48 that are capable of coupling directly to the fiber optic link 40. For  
15 of directly coupling to the fiber optic cable 40, a general purpose media processor 12 can be used as an interface to other forms of network interconnection.

As shown in FIG. 2, alternate forms of interconnection consist of coaxial cable lines 42 and twisted pair wiring 44. Coaxial cable lines are currently in place throughout the U.S. and is typically employed to provide cable television  
20 services to residential homes. According to the preferred embodiment of the invention, general purpose media processors 12 can be installed at these residential locations 52. In contrast to the specialized processor approach, the general purpose media processor 12 provides enough bandwidth to allow for bi-directional communications to and from these residential locations 52.

25 Network servers 46 controlled by general purpose media processors 12 are also employed throughout the network 38. For example, the network servers 46 can be used to interface between the fiber optic network 40 and twisted pair wiring 44. Twisted pair wiring 44 is still employed for small businesses 50 and residential locations 52 that do not or cannot currently subscribe to coaxial  
30 cable or fiber optic network services. General purpose media processors 12 are also disposed at these small business locations 50 and non-cable residential

- 15 -  
16

locations 52. General purpose media processors 12 are also installed in wireless or mobile locations 52, which are coupled to the network 38 through agile digital radios (not shown). As shown in FIG. 2, network databases or other peripherals 56 can also coupled to general purpose media processors 12 in the network 38.

5           The general purpose media processor 12 is operable at significantly high bandwidths in order to receive, process and transmit unified media data streams. Referring to FIG. 3, the respective frequencies for various types of media data streams are set forth against a bandwidth spectrum 60. The bandwidth spectrum 60 includes three component spectrums, all along the same range of  
10       frequencies, which represent the various frequency rates of digital media communications. Current computing bandwidth capabilities are also displayed. The telecommunications spectrum 62 shows the various frequency bands used for telecommunications transmission. For example, teletype terminals and modems operate in a range between approximately 64 bits/second to 16 kilobits/second.  
15       The ISDN telecommunication protocol operates at 64 kilobits/second. At the upper end of the telecommunications spectrum 62, T1 and T3 trunks operate at one megabit per second and 32 megabits per second, respectively. The SONET frequency range extends from approximately 128 megabits per second up to approximately 32 gigabits per second. Accordingly, in order to carry such broad  
20       band communications, the general purpose media processor 12 is capable of transferring information at rates into the gigabits per second range or higher.

          A spectrum of typical media data streams is presented in the media spectrum 64 shown in FIG. 3. Voice and music transmissions are centered at frequencies of approximately 64 kilobits per second and one megabit per second,  
25       respectively. At the upper end of the media spectrum 64, video transmission takes place in a range from 128 megabits per second for high density television up to over 256 gigabits per second for movie applications. When using common video compression techniques, however, the video transmission spectrum can be shifted down to between 32 kilobits per second to 128 megabits per second as a result of  
30       the data compression. As described below, the processing required to achieve the data compression results in an increase in bandwidth requirements.



Current computing bandwidths are shown in the computing spectrum 66 of FIG. 3. Serial communications presently take place in a range between two kilobits per second up to 512 kilobits per second. The Ethernet network protocol operates at approximately 8 megabits per second. Current dynamic random access memory and other digital input/output peripherals operate between 32 megabits per second and 512 megabits per second. Presently available microprocessors are capable of operation in the low gigabits per second range. For example, the '386 Pentium microprocessor manufactured by Intel Corporation of Santa Clara, California operates in the lower half of that range, and the Alpha microprocessor manufactured by Digital Equipment Corporation approaches the 16 gigabits per second range.

When video compression is employed, as expressed above, the associated processing overhead reduces the effective bandwidth of the particular processor. As a result, in order to handle compressed video, these processors must operate in the terahertz frequency range. The bandwidth spectrum 60 shown in FIG. 4 represents the effect of handling media data streams including compressed video. The computing spectrum 66 is skewed down to properly align the computing bandwidth requirements with the telecommunications spectrum 62 and the media spectrum 64. Accordingly, current processor technology is not sufficient to handle the transmission and processing associated with complex streams of multi-media data.

The current specialized processor approach to media processing is illustrated in the block diagram shown in FIG. 5. As shown in FIG. 5, special purpose processors are coupled to a back plane 70, which is capable of transmitting instructions and data at the upper kilobits to lower gigabits per second range. In a typical configuration, an audio processor 76, video processor 78, graphics processor 80 and network processor 82 are all coupled to the back plane 70. Each of the audio, video, graphics and network processors 76-82 typically employ their own private or dedicated memories 84, which are only accessible to the specific processor and not accessible over the back plane 70. As described above, however, unless video data streams are constantly being processed, for

example, the video processor 78 will sit idle for periods of time. The computing power of the dedicated video processor 78 is thus only available to handle video data streams and is not available to handle other media data streams that are directed to other dedicated processors. This, of course, is an inefficient use of the video processor 78 particularly in view of the overall processing capability of this multi-processor system.

The general purpose media processor 12, in contrast, handles a data stream of audio, video, graphics and network information all at the same time with the same processor. In order to handle the ever changing combination of data types, the general purpose media processor 12 is dynamically partitionable to allocate the appropriate amount of processing for each combination of media in a unified media data stream. A block diagram of two preferred general purpose media processor system configurations is shown in FIG. 6. Referring to FIG. 6(a), a general purpose media processor 12 is coupled to a high-speed back plane 90. The presently preferred back plane 90 is capable of operation at 30 gigabits per second. As those skilled in the art will appreciate, back planes 90 that are capable of operation at 400 gigabits per second or greater bandwidth are envisioned within the spirit and scope of the invention. Multiple memory devices 92 are also coupled to the back plane 90, which are accessible by the general purpose media processor 12. Input/output devices 94 are coupled to the back plane 90 through a dual-ported memory 92. The configuration of the input/output devices 94 on one end of the dual-ported memory 92 allows the sharing of these memory devices 92 throughout a network 38 of general purpose media processors 12.

Alternatively, FIG. 6(b) shows a presently preferred integrated general purpose media processor 12. The integrated processor includes on-board memory and I/O 86. The on-board memory is preferably of sufficient size to optimize throughput, and can comprise a cache and/or buffer memory or the like. The integrated media processor 12 also connects to external memory 88, which is preferably larger than the on-board memory 86 and forms the system main memory.

### Execution Unit

One presently preferred embodiment of an integrated general purpose media processor 12 is shown in FIG. 7. The core of the integrated general purpose media processor 12 comprises an execution unit 100. Three main  
5 elements or subsections are included in the execution unit 100. A multiple precision arithmetic/logic unit ("ALU") 102 performs all logical and simple arithmetic operations on incoming media data streams. Such operations consist of calculate and control operations such as Boolean functions, as well as addition, subtraction, multiplication and division. These operations are performed on single  
10 or unified media data streams transmitted to and from the multiple precision ALU 102 over a data bus or data path 108. Preferably the data path 108 is 128 bits wide, although those skilled in the art will appreciate that the data path 108 can take on any width or size without departing from the spirit and scope of the invention. The wider the data path 108 the more unified media data can be  
15 processed in parallel by the general purpose media processor 12.

Coupled to the multi-precision ALU 102 via the data path 108, and also an element of the execution unit 100, is a programmable switch 104. The programmable switch 104 performs data handling operations on single or unified media data streams transmitted over the data path 108. Examples of such data  
20 handling operations include deals, shuffles, shifts, expands, compresses, swizzles, permutes and reverses, although other data handling operations are contemplated. These operations can be performed on single bits or bit fields consisting of two or more bits up to the entire width of the data path 108. Thus, single bits or bit fields of various sizes can be manipulated through programmable operation of the  
25 switch 104.

Examples of the presently preferred data manipulation operations performed by the general purpose media processor 12 are shown in FIG. 8. A group expand operation is visually illustrated in FIG. 8(a). According to the group expand operation, a sequential field of bits 270 can be divided into  
30 constituent sub-fields 272a-272d for insertion into a larger field array 274. The reverse of the group expand operation is a group compress or extract operation. A

19  
20

visual illustration of the group compress or extract operation is shown in FIG. 8(b). As shown, separate sub-fields 272a-272d from a larger bit field 274 can be combined to form a contiguous or sequential field of bits 270.

Referring to FIGS. 8(c)-8(e), group deal, shuffle, swizzle and  
5 permute operations performed by the programmable switch 104 are also  
illustrated. The operations performed by these instructions are readily understood  
from a review of the drawings. The group manipulation operations illustrated in  
FIGS. 8(a)-8(e) comprise the presently contemplated data manipulation operations  
for the general purpose media processor 12. As those skilled in the art will  
10 appreciate, either a subset of these operations or additional data manipulation  
operations can be incorporated in other alternate embodiments of the general  
purpose media processor 12 without departing from the spirit and scope of the  
invention.

Referring again to FIG. 7, higher level mathematical operations than  
15 those performed by the multi-precision ALU 102 are performed in the general  
purpose media processor 12 through an extended math element 106. The extended  
math element 106 is coupled to the data path 108 and also comprises part of the  
execution unit 100. The extended math element 106 performs the complex  
arithmetic operations necessary for video data compression and similarly intensive  
20 mathematical operations. One presently preferred example of an extended math  
operation comprises a Galois field operation. Other examples of extended  
mathematical functions performed by the extended math element 106 include CRC  
generation and checking, Reed-Solomon code generation and checking, and  
spread-spectrum encoding and decoding. As those skilled in the art appreciate,  
25 additional mathematical operations are possible and contemplated.

According to the preferred embodiment of the integrated general  
purpose media processor 12, a register file 110 is provided in addition to the  
execution unit 100 to process media data. The register file 110 stores and  
transmits data streams to and from the execution unit 100 via the data path 108.  
30 Rather than employing a complex set of specific or dedicated registers, the general  
purpose media processor 12 preferably includes 64 general purpose registers in the

register file 110 along with one program counter (not shown). The 64 general purpose registers contained in the register file 110 are all available to the user/programmer, and comprise a portion of the user state of the general purpose media processor 12. The general purpose registers are preferably capable of storing any form of data. Each register within the register file 110 is coupled to the data path 108 and is accessible to the execution unit 100 in the same manner. Thus, the user can employ a general purpose register according to the specific needs of a particular program or unique application. As those skilled in the art will appreciate, the register file 110 can also comprise a plurality of register files 110 configured in parallel in order to support parallel multi-threaded processing.

#### Instruction Set and User Programming

Control or manipulation of data processed by the general purpose media processor 12 is achieved by selected instructions programmed by the user. Those skilled in the art will appreciate that a great number of programs are possible through various sequences of instructions. Particular programs can be developed for each unique implementation of the general purpose media processor 12. A detailed discussion of such specific programs is therefore beyond the scope of this description.

One presently preferred instruction set for the general purpose media processor 12 is included in the Microfiche Appendix, the contents of which are hereby incorporated herein by reference. A list of the presently preferred major operation codes for the general purpose media processor 12 appears below in Table I.

21



# MAJOR OPERATION CODES

MAJOR	0	32	64	96	128	160	192	224
0	ERES	GSHUFFLEI	FMULADD16	GMULADD1	LU16LAI	SAAS64LAI	EADDIO	BFE16
1	ESHUFFLEI4MUX	GSHUFFLEI4MUX	FMULADD32	GMULADD2	LU16BAI	SAAS64BAI	EADDIUO	BFNUE16
2		GSELECT8	FMULADD64	GMULADD4	LU16LI	SCAS64LAI	ESETIL	BFNUGE16
3	EMDEPI	GMDEPI		GMULADD8	LU16BI	SCAS64BAI	ESETIGE	BFNUL16
4	EMUX	GMUX	FMULSUB16	GMULADD16	LU32LAI	SMAS64LAI	ESETIE	BFE32
5	E8MUX	G8MUX	FMULSUB32	GMULADD32	LU32BAI	SMAS64BAI	ESETINE	BFNUE32
6	EGFMUL64	GGFMUL8	FMULSUB64	GMULADD64	LU32LI	SMUX64LAI	ESETIUL	BFNUGE32
7	ETRANSPOSE8MUX	GTRANSPOSE8MUX		GEXTRACT128	LU32BI	SMUX64BAI	ESETIUGE	BFNUL32
8					L16LAI	S16LAI	ESUBIO	BFE64
9	ESWIZZLE	GSWIZZLE		GUMULADD2	L16BAI	S16BAI	ESUBIUO	BFNUE64
10		GSWIZZLECOPY		GUMULADD4	L16LI	S16LI	ESUBIL	BFNUGE64
11		GSWIZZLESWAP		GUMULADD8	L16BI	S16BI	ESUBIGE	BFNUL64
12	EDEPI	GDEPI	F.16	GUMULADD16	L32LAI	S32LAI	ESUBIE	BFE128
13	EUDEPI	GUDEPI	F.32	GUMULADD32	L32BAI	S32BAI	ESUBINE	BFNUE128
14	EWTHI	GWTHI	F.64	GUMULADD64	L32LI	S32LI	ESUBIUL	BFNUGE128
15	EUWTHI	GUWTHI		GUEXTRACT128	L32BI	S32BI	ESUBIUGE	BFNUL128
16			GFMULADD16	GEXTRACTI	L64LAI	S64LAI	EADDI	BANDE
17			GFMULADD32	GEXTRACTI16	L64BAI	S64BAI	EXORI	BANDNE
18			GFMULADD64	GEXTRACTI32	L64LI	S64LI	EORI	BL/BLZ
19			GFMULADD128	GUEXTRACTI64	L64BI	S64BI	EANDI	BGE/BGEZ
20			GFMULSUB16	GEXTRACT	L128LAI	S128LAI	ESUBI	BE
21			GFMULSUB32	.I.64	L128BAI	S128BAI		BNE
22			GFMULSUB64	G.EXTRACT	L128LI	S128LI	ENORI	BUL/BGZ
23			GFMULSUB128	.I.128	L128BI	S128BI	ENANDI	BUGE/BLEZ
24				G.1	L8I	S8I		BGATEI
25				G.2	LU8I			
26				G.4				
27				G.8				
28		ECOPYI	GF.16	G.16			ECOPYI	BI
29			GF.32	G.32				BLINKI
30			GF.64	G.64				
31		E.MINOR	GF.128	G.128	L.MINOR	S.MINOR	E.MINOR	B.MINOR

major operation code field values

TABLE I

As shown in Table I, the major operation codes are grouped according to the function performed by the operations. The operations are thus arranged and listed above according to the presently preferred operation code number for each instruction. As many as 255 separate operations are contemplated for the preferred embodiment of the general purpose media processor 12. As shown in Table I, however, not all of the operation codes are presently implemented. As those skilled in the art will appreciate, alternate schemes for organizing the operation codes, as well as additional operation codes for the general purpose media processor 12, are possible.

22-  
B

The instructions provided in the instruction set for the general purpose media processor 12 control the transfer, processing and manipulation of data streams between the register file 110 and the execution unit 100. The presently preferred width of the instruction path 112 is 32-bits wide, organized as four eight-bit bytes ("quadlets"). Those skilled in the art will appreciate, however, that the instruction path 112 can take on any width without departing from the spirit and scope of the invention. Preferably, each instruction within the instruction set is stored or organized in memory on four-byte boundaries. The presently preferred format for instructions is shown in FIG. 9(a).

As shown in FIG. 9(a), each of the presently preferred instruction formats for the general purpose media processor 12 includes a field 280 for the major operation code number shown in Table I. Based on the type of operation performed, the remaining bits can provide additional operands according to the type of addressing employed with the operation. For example, the remainder of the 32-bit instruction field can comprise an immediate operand ("imm"), or operands stored in any of the general registers ("ra," "rb," "rc," and "rd"). In addition, minor operation codes 282 can also be included among the operands of certain 32-bit instruction formats.

The presently preferred embodiment of the general purpose media processor 12 includes a limited instruction set similar to those seen in Reduced Instruction Set Computer ("RISC") systems. The preferred instruction set for the general purpose media processor 12 shown in Table I includes operations which implement load, store, synchronize, branch and gateway functions. These five groups of operations can be visually represented as two general classes of related operations. The branch and gateway operations perform related functions on media data streams and are thus visually represented as block 114 in FIG. 7. Similarly, the load, store and synchronize operations are grouped together in block 116 and perform similar operations on the media data streams. (Blocks 114 and 116 only represent the above classification of these operations and their function in the processing of media data streams, and do not indicate any specific underlying electronic connections.) A more detailed discussion of these operations, and the

23  
24


functionality of the general purpose media processor 12, appears in the Microfiche Appendix.

5 The four-byte structure of instructions for the general purpose media processor 12 is preferably independent of the byte ordering used for any data structures. Nevertheless, the gateway instructions are specifically defined as 16-byte structures containing a code address used to securely invoke a procedure at a higher privilege level. Gateways are preferably marked by protection information specified in the translation lookaside buffer 148 in the memory management unit 122. Gateways are thus preferably aligned on 16-byte boundaries in the external  
10 memory. In addition to the general purpose registers and program counter, a privilege level register is provided within the register file 110 that contains the privilege level of the currently executing instruction.

The instruction set preferably includes load and store instructions that move data between memory and the register file 110, branch instructions to  
15 compare the content of registers and transfer control, and arithmetic operations to perform computations on the contents of registers. Swap instructions provide multi-thread and multi-processor synchronization. These operations are preferably indivisible and include such instructions as add-and-swap, compare-and-swap, and multiplex-and-swap instructions. The fixed-point compare-and-branch instructions  
20 within the instruction set shown in Table I provide the necessary arithmetic tests for equality and inequality of signed and unsigned fixed-point values. The branch through gateway instruction provides a secure means to access code at a higher privileged level in a form similar to a high level language procedure call generally known in the art.

25 The general purpose media processor 12 also preferably supports floating-point compare-and-branch instructions. The arithmetic operations, which are supported in hardware, include floating-point addition, subtraction, multiplication, division and square root. The general purpose media processor 12 preferably supports other floating-point operations defined by the ANSI-IEEE  
30 floating-point standard through the use of software libraries. A floating point

- 24 -






value can preferably be 16, 32, 64 or 128-bits wide. Examples of the presenting preferred floating-point data sizes are illustrated in FIG. 9(b).

The general purpose media processor 12 preferably supports virtual memory addressing and virtual machine operation through a memory management unit 122. Referring to FIG. 10(a), one presently preferred embodiment of the  
5 memory management unit 122 is shown. The memory management unit 122 preferably translates global virtual addresses into physical addresses by software programmable routines augmented by a hardware translation lookaside buffer ("TLB") 148. A facility for local virtual address translation 164 is also preferably  
10 provided. As those skilled in the art will appreciate, the memory management unit 122 includes a data cache 166 and a tag cache 168 that store data and tags associated with memory sections for each entry in the TLB 148.

A block diagram of one preferred embodiment of the TLB 148 is shown in FIG. 10(b). The TLB 148 receives a virtual address 230 as its input.  
15 For each entry in the TLB 148, the virtual address 230 is logically AND-ed with a mask 232. The output of each respective AND gate 234 is compared via a comparator 236 with each entry in the TLB 148. If a match is detected, an output from the comparator 236 is used to gate data 240 through a transceiver 238. As those skilled in the art will appreciate, a match indicates the entry of the  
20 corresponding physical address within the contents of the TLB 148 and no external memory or I/O access is required. The data 240 for the data cache 166 (FIG. 10(a)) is then combined with the remaining lower bits of the virtual address 230 through an exclusive-OR gate 242. The resultant combination is the physical address 244 output from the TLB 148. If a match is not detected between the  
25 logical address and the contents of the tag cache 168, the memory management unit 122 an external memory or I/O access is necessary to retrieve the relevant portion of memory and update the contents of the TLB 148 accordingly.

Using generally known memory management techniques, the memory management unit 122 ensures that instructions (and data) are properly  
30 retrieved from external memory (or other sources) over an external input/output bus 126 (see FIG. 7). As described in more detail below, a high bandwidth

- 25 -  


interface 124 is coupled to the external input/output bus 126 to communicate instructions (and media data streams) to the general purpose media processor 12. The presently preferred physical address width for the general purpose media processor 12 is eight bytes (64-bits). In addition, the memory management unit  
5 122 preferably provides match bits (not shown) that allow large memory regions to be assigned a single TLB entry allowing for fine grain memory management of large memory sections. The memory management unit 122 also preferably includes a priority bit (not shown) that allows for preferential queuing of memory areas according to respective levels of priority. Other memory management  
10 operations generally known in the art are also performed by the memory management unit 122.

Referring again to FIG. 7, instructions received by the general purpose media processor 12 are stored in a combined instruction buffer/cache 118. The instruction buffer/cache 118 is dynamically subdivided to store the largest  
15 sequence of instructions capable of execution by the execution unit 100 without the necessity of accessing external memory. In a preferred embodiment of the invention, instruction buffer space is allocated to the smallest and most frequently executed blocks of media instructions. The instruction buffer thus helps maintain the high bandwidth capacity of the general purpose media processor 12 by  
20 sustaining the number of instructions executed per second at or near peak operation. That portion of the instruction buffer/cache 118 not used as a buffer is, therefore, available to be used as cache memory. The instruction buffer/cache 118 is coupled to the instruction path 112 and is preferably 32 kilobytes in size.

A data buffer/cache 120 is also provided to store data transmitted  
25 and received to and from the execution unit 100 and register file 110. The data buffer/cache 120 is also dynamically subdivided in a manner similar to that of the instruction buffer/cache 118. The buffer portion of the data buffer/cache 120 is optimized to store a set size of unified media data capable of execution without the necessity of accessing external memory. In a preferred embodiment of the  
30 invention, data buffer space is allocated to the smallest and most frequently accessed working sets of media data. Like the instruction buffer, the data buffer

thus maintains peak bandwidth of the general purpose media processor 12. The data buffer/cache 120 is coupled to the data path 108 and is preferably also 32 kilobytes in size.

5 The preferred embodiment of the general purpose media processor 12 includes a pipelined instruction pre-fetch structure. Although pipelined operation is supported, the general purpose media processor 12 also allows for non-pipelined operations to execute without any operational penalty. One preferred pipeline structure for the general purpose media processor 12 comprises a "super-string" pipeline shown in FIG. 11. A super-string pipeline is designed to  
10 fetch and execute several instructions in each clock cycle. The instructions available for the general purpose media processor 12 can be broken down into five basic steps of operation. These steps include a register-to-register address calculation, a memory load, a register-to-register data calculation, a memory store and a branch operation. According to the super-string pipeline organization of the  
15 general purpose media processor 12, one instruction from each of these five types may be issued in each clock cycle. The presently preferred ordering of these operations are as listed above where each of the five steps are assigned letters "A," "L," "E," "S" and "B" (see FIG. 11).

According to the super-string pipelining technique, each of the  
20 instructions are serially dependent, as shown in FIG. 11, and the general purpose media processor 12 has the ability to issue a string of dependent instructions in a single clock cycle. These instructions shown in FIG. 11 can take from two to five cycles of latency to execute, and a branch prediction mechanism is preferably used to keep up the pipeline filled (described below). Instructions can be encoded in  
25 unit categories such as address, load, store/sync, fixed, float and branch to allow for easy decoding. A similar scheme is employed to pre-fetch data for the general purpose media processor 12.

As those skilled in the art will appreciate, the super-string pipeline can be implemented in a multi-threaded environment. In such an implementation,  
30 the number of threads is preferably relatively prime with respect to functional unit

28


rates so that functional units can be scheduled in a non-interfering fashion between each thread.

In another more preferred embodiment, a "super-spring" pipelining scheme is employed with the general purpose media processor 12. The super-spring pipeline technique breaks the super-string pipeline shown in FIG. 11 into two sections that are coupled via a memory buffer (not shown). A visual representation of the super-spring pipeline technique is shown in FIG. 12. The front of the pipeline 204, in which address calculation (A), memory load (L), and branch (B) operations are handled, is decoupled from the back of the pipeline 206, in which data calculation (E) and memory store (S) operations are handled. The decoupling is accomplished through the memory buffer (not shown), which is preferably organized in a first-in-first-out ("FIFO") fast/dense structure. (The memory buffer is functionally represented as a spring in FIG. 12.)

As indicated in Table I above, the general purpose media processor 12 does not include delayed branch instructions, and so relies upon branch or fetch prediction techniques to keep the pipeline full in program flows around unconditional and conditional branch instructions. Many such techniques are generally known in the art. Examples of some presently preferred techniques include the use of group compare and set, and multiplex operations to eliminate unpredictable branches; the use of short forward branches, which cause pipeline neutralization; and where branch and link predicts the return address in a one or more entry stack. In addition, the specialized gateway instructions included in the general purpose media processor 12 allow for branches to and from protected virtual memory space. The gateway instructions, therefore, allow an efficient means to transfer between various levels of privilege.

As described above, two basic forms of media data are processed by the general purpose media processor 12, as shown in FIG. 7. These data streams generally comprise Nyquist sampled I/O 128, and standard memory and I/O 130. As shown in FIG. 7, audio 132, video 134, radio 136, network 138, tape 140 and disc 142 data streams comprise some examples of digitally sampled I/O 128. As those skilled in the art will appreciate, other forms of digitally sampled I/O are

- 28 -



contemplated for processing by the general purpose media processor 12 without departing from the spirit and scope of the invention. Standard memory and I/O 130 comprises data received and transmitted to and from general digital peripheral devices used in the design of most computer systems. As shown in FIG. 7, some  
5 examples of such devices include dynamic random access memory ("DRAM") 146, or any data received over the PCI bus 144 generally known in the art. Other forms of standard memory and I/O sources are also contemplated. The various fixed-point data sizes preferred for the general purpose media processor 12 are illustrated in FIG. 9(c).

#### External Interface

As mentioned above, the general purpose media processor 12 includes a high bandwidth interface 124 to communicate with external memory and input/output sources. As part of the high bandwidth interface 124, the general  
15 purpose media processor 12 integrates several fast communication channels 156 (FIG. 13) to communicate externally. These fast communication channels 156 preferably couple to external caches 150, which serve as a buffer to memory interfaces 152 coupled to standard memory 154. The caches 150 preferably comprise synchronous static random access memory ("SRAM"), each of which are  
20 sixty-four kilobytes in size; and the standard memories 154 comprise DRAM's. The memory interfaces 152 transmit data between the caches 150 and the standard memories 154. The standard memories 154 together form the main external memory for the general purpose media processor 12. The cache 150, memory interface 152, standard memory 154 and input/output channel 156 therefore make  
25 up a single external memory unit 158 for the general purpose media processor 12.

According to the presently preferred embodiment of the invention, the memory interface protocol embeds read and write operations to a single memory space into packets containing command, address, data and acknowledgment information. The packets preferably include check codes that  
30 will detect single-bit transmission errors and some multiple-bit errors. As many as eight operations may be in progress at a time in each external memory unit 158.

As shown in FIG. 13, up to four external memory units 158 may be cascaded together to expand the memory available to the general purpose media processor 12, and to improve the bandwidth of the external memory. Through such cascaded memory units 158, the memory interface 152 provides for the direct  
5 connection of multiple banks of standard memory 154 to maintain operation of the general purpose media processor 12 at sustained peak bandwidths.

According to one embodiment shown in FIG. 13, up to four standard memory devices 154 can be coupled to each memory interface 152. Each standard memory 154 thus includes as many as four banks of DRAM, each of  
10 which is preferably sixteen bits wide. The standard memories 154 are connected in parallel to the memory interface 152 forming a 72-bit wide data bus 160, where 64 bits are preferably provided for data transfer and eight bits are provided for error correction. In addition to the data bus 160, an address/control bus 162 is coupled between the memory interface 152 and each standard memory 154. The  
15 address/control bus 162 preferably comprises at least twelve address lines (4 kilobits x 16 memory size) and four control lines as shown in FIG. 13. An alternate manner for coupling the DRAM's to the memory interface 152 is illustrated in FIG. 14. As shown in FIG. 14, two banks of four DRAM single in-line memory modules are coupled in parallel to the memory interface 152. The  
20 memory interface 152 also supports interleaving to enhance bandwidth, and page mode accesses to improve latency for localized addressing.

Using standard DRAM components, the external memory units 158 achieve bandwidths of approximately two gigabits/second with the standard memories 154. When four such external memory units 158 are coupled via the  
25 communication channel 156, therefore, the total bandwidth of the external main memory system increases to one gigabyte/second. As discussed further below, in implementations with two or eight communication channels 156, the aggregate bandwidth increases to two and eight gigabytes/second, respectively.

A more detailed depiction of the communication channel 156  
30 circuitry appears in FIG. 15. According to the preferred embodiment of the invention, each communication channel 156 comprises two unidirectional, byte-

wide, differential, packet-oriented data channels 156a, 156b (see FIG. 13). As explained above, where memory units 158 are cascaded together in series, the output of one memory unit 158 is connected to the input of another memory unit 158. The two unidirectional channels are thus connected through the memory units 158 forming a loop structure and make up a single bi-directional memory interface channel.

Referring to FIG. 15, each communication channel 156 is preferably eight bits wide, and each bit is transmitted differentially. For example, output transceiver 170 for bit  $D_{0out}$  transmits both  $D_0$  and  $/D_0$  signals over the communication channel 156. Additional transceivers are similarly provided for the remaining bits in the channel 156. (The transceiver 176 for bit  $D_{7out}$  and associated differential lines 178, 180 are shown in FIG. 15.) A  $CLK_{out}$  transceiver 182 is also provided to generate differential clock outputs 184, 186 over the channel 156. To complete the link between memory units 158, input transceivers 188-192 are provided in each memory unit 158 for each of the differential bits and clock signals transmitted over the communication channel 156. These input signals 172, 174, 178, 180, 184, 186 are preferably transmitted through input buffers 194-198 to other parts of the memory unit 158 (described above).

Each memory unit 158 also includes a skew calibrator 200 and phase locked loop ("PLL") 202. The skew calibrator 200 is used to control skew in signals output to the communication channel 156. Preferably, digital skew fields are employed, which include set numbers of delay stages to be inserted in the output path of the communication channel 156. Setting these fields, and the corresponding analog skew fields, permits a fine level of control over the relative skew between output channel signals.

The PLL 202 recovers the clock signal on either side of the communication channel 156 and is thus provided to remove clock jitter. The clock signals 184, 186 preferably comprise a single phase, constant rate clock signal. The clock signals 184, 186 thus contain alternating zero and one values transmitted with the same timing as the data signals 172, 174, 178, 180. The clock signal frequency is, therefore, one-half the byte data rate. The communication channel

31  
32

156 preferably operates at constant frequency and contains no auxiliary control, handshaking or flow control information.

Each external memory unit 158 preferably defines two functional regions: a memory region, implemented by the cache 150 backed by standard memory 154 (see FIG. 13), and a configuration region, implemented by registers (not shown). Both regions are accessed by separate interfaces; the communication channel 156 is used to access the memory region, and a serial interface (described below) is used to access the configuration region. In the memory region, the caches 150 are preferably write-back (write-in) single-set (direct-map) caches for data originally contained in standard memory 154. All accesses to memory space should maintain consistency between the contents of the cache 150 and the contents of the standard memory 154. The configuration region registers provide the mechanism to detect and adjust skew in the communication channel 156. Software is preferably employed to adaptively adjust the skew in the channel 156 through digital skew fields, as explained above. The serial interface thus is used to configure the external memory units 158, set diagnostic modes and read diagnostic information, and to enable the use of a high-speed tester (not shown).

One presently preferred embodiment of the invention employs two byte-wide packet communication channels 156 (FIG. 16(a)). In order to further increase the bandwidth of the general purpose media processor 12, up to sixteen byte-wide packet communication channels 156 can be employed. Referring to FIG. 16(b), twelve communication channels, comprising eight memory channels 210, a ninth channel for parallel processing 212 (described below), and three input/output ("I/O") channels 214, are shown. Each of the communication channels 210-214 preferably employs the cascade configuration of four channel interface devices 216. (Each channel interface device 216 coupled to the memory channels 210 corresponds to the external memory unit 158 shown in FIG. 13.) Through each of the twelve communication channels shown in FIG. 16(b), the general purpose media processor 12 can request or issue read or write transactions. When not interleaved, the twelve channels provide a single contiguous memory space for each channel interface device 216.



Alternatively, memory accesses may be interleaved in order to provide for continuous access to the external memory system at the maximum bandwidth for the DRAM memories. In an interleaved configuration, at any point in time some memory devices will be engaged in row pre-charge, while others may be driving or receiving data, or receiving row or column addresses. The memory interface 152 (FIG. 13) thus preferably maps between a contiguous address space and each of the separate address spaces made available within each external memory unit 158. For maximum performance, therefore, the memory interface is interleaved so that references to adjacent addresses are handled by different memory devices. Moreover, in the preferred embodiment, additional memory operations may be requested before the corresponding DRAM bank is available. In an interleaved approach, these operations are placed in a queue until they can be processed. According to the preferred embodiment, memory writes have lower priority than memory reads, unless an attempt is made to read an address that is queued for a write operation. As those skilled in the art will appreciate, the depth of the memory write queue is dictated by the specific implementation.

Although up to four external memory units 158 are preferably cascaded to form effectively larger memories, some amount of latency may be introduced by the cascade. Packets of data transmitted over the communication channel 156 are uniquely addressed to a particular channel interface device 216. A packet received at a particular device, which specifies another module address, is automatically passed to the correct channel interface device 216. Unless the module address matches a particular device 216, that packet simply passes from the input to the output of the interface device 216. This mechanism divides the serial interconnection of interface devices 216 into strings, which function as a single larger memory or peripheral, but with possibly longer response latency.

In addition to the memory channels 210, the general purpose media processor 12 provides several communication channels 214 for communication with external input/output devices. Referring to FIG. 16(b), three input/output channels 214 having SRAM buffered memory (see FIG. 13) provide an interface

33  
34

to external standard I/O devices (not shown). Like the eight memory channels 210, the three I/O channels 214 are byte-wide input/output channels intended to operate at rates of at least one gigahertz. The three I/O channels 214 also operate as a packet communication link to synchronous SRAM memory 208 within the channel interface device 216. A controller 226 within the channel interface device 216 completes the interface to the I/O devices.


The three I/O channels 214 preferably function in like manner to the memory channels 210 described above. The interface protocol for the three I/O channels 214 divides read and write operations to a single memory space into packets containing command, address, data and acknowledgment information. The packets also include a check code that will detect single-bit transmission errors and some multiple-bit errors. According to the preferred embodiment of the invention, as many as eight operations may progress in each interface device 216 at a time. As shown in FIG. 16(b), up to four channel interface devices 216 can be cascaded together to expand the bandwidth in the three I/O channels 214. A bit-serial interface (not shown) is also provided to each of the channel interface devices 216 to allow access to configuration, diagnostic and tester information at standard TTL signal levels at a more moderate data rate. (A more detailed description of the serial interface is provided below).

Like the memory channels 210, each I/O channel 214 includes nine signals -- one clock signal and eight data signals. Differential voltage levels are preferably employed for each signal. Each channel interface device 216 is preferably terminated in a nominal 50 ohm impedance to ground. This impedance applies for both inputs and outputs to the communication channel 156. A programmable termination impedance is preferred.

#### Interface Communication

According to one presently preferred embodiment of the invention, the channel interface devices 216 can operate as either master devices or slave devices. A master device is capable of generating a request on the communication channel 156 and receiving responses from the communication channel 156. Slave

34 -



devices are capable of receiving requests and generating responses, over the communication channel 156. A master device is preferably capable of generating a constant frequency clock signal and accepting signals at the same clock frequency over the communication channel 156. A slave device, therefore, should  
5 operate at the same clock rate as the communication channel 156, and generate no more than a specified amount of variation in output clock phase relative to input clock phase. The master device, however, can accept an arbitrary input clock phase and tolerates a specified amount of variation in clock phase over operating conditions.

10 Packets of information sent over the communication channel 156 preferably contain control commands, such as read or write operations, along with addresses and associated data. Other commands are provided to indicate error conditions and responses to the above commands. When the communication channel 156 is idle, such as during initialization and between transmitted packets,  
15 an idle packet, consisting of an all-zero byte and an all- one byte is transmitted through the communication channel 156. Each non-idle packet consists of two bytes or a multiple of two bytes, and begins with a byte having a value other than all zeros. All packets transmitted over the communication channel 156 also begin during a clock period in which the clock signal is zero, and all packets preferably  
20 end during a clock period in which the clock signal is one. A depiction of the preferred packet protocol format for transmission over the communication channel 156 appears in FIG. 17.

The general form of each packet is an array of bytes preferably without a specific byte ordering. The first byte contains a module address  
25 ("ma") in the high order two bits; a packet identifier, usually a command 252 ("com"), in the next three bit positions; and a link identification number 254 ("lid") in the last three bit positions. The interpretation of the remaining bytes of a packet depend upon the contents of the packet identifier. The length of each packet is preferably implied by the command specified in the initial byte of the  
30 packet. A check byte is provided and computed as odd bit-wise parity with a leftward circular rotation after accumulating each byte. This technique provides

detection of all single-bit and some multiple-bit errors, but no correction is provided.

5           The modular address 250 field of each packet is preferably a two-bit field and allows for as many as four slave devices to be operated from a single communication channel 156. Module address values can be assigned in one of two fashions: either dynamically assigned through a configuration register (not shown), or assigned via static/geometric configuration pins. Dynamic assignment through a configuration register is the presently preferred method for assigning module address values.

10           The link identification number 254 field is preferably 3-bits wide and provides the opportunity for master devices to initiate as many as eight independent operations at any one time to each slave device. Each outstanding operation requires a distinct link identification number, but no ordering of operations should be implied by the value of the link identification field. Thus,  
15           there is preferably no requirement for link identification values 254 to be sequentially assigned either in requests or responses.

          The receipt of packets over the communication channel 156 that do not conform to the channel protocol preferably generates an error condition. As those skilled in the art will appreciate, the level or degrees to which a specific  
20           implementation detects errors is defined by the user. In one presently preferred embodiment of the invention, all errors are detected, and the following protocol is employed for handling errors. For each error detected, the channel interface device 216 causes a response explicitly indicating the error condition. Channel interface devices 216 reporting an invalid packet will then suppress the receipt of  
25           additional packets until the error is cleared. The transmitted packet is otherwise ignored. However, even though the erroneous packet is ignored, the channel interface devices 216 preferably continue to process valid packets that have already been received and generate responses thereto. An identification of the presently preferred commands 252 to be used over the communication channel 156 are listed  
30           in FIG. 17.

In the master/slave preferred embodiment, the channel interface devices 216 forward packets that are intended for other devices connected to the communication channel 156, as described above. In slave devices, forwarding is performed based on the module address 250 field of the packet. Packets which  
5 contain a module address 250 other than that of the current device are forwarded on to the next device. All non-idle packets are thus forwarded including error packets. In master devices, forwarding is performed based on the link identifier number 254 of the packet. Packets that contain link identifier numbers 254 not generated by the specific channel interface device 216 are forwarded. In order to  
10 reduce transmission latency, a packet buffer may be provided. As those skilled in the art appreciate, the suitable size for the packet buffer depends on the amount of latency tolerable in a particular implementation.

A variety of master/slave ring configurations are possible using the high bandwidth interface 124 of the invention. Five ring configurations are  
15 currently preferred: single-master, dual-master, multiple-master, single-slave and multiple-master/multiple-slave. The simplest ring configuration contains a single non-forwarding master device and a single non-forwarding slave device. No forwarding is required for either device in this configuration as packets are sent directly to the recipient. A single-master ring, however, may contain a cascade of  
20 up to four slave devices (see FIGS. 13, 16). In the single-master ring configuration, each slave device is configured to a distinct module address, and each slave device forwards packets that contain module address fields unequal to their own. As discussed above, a single-master ring provides a larger memory or I/O capacity than a master-slave pair, but also introduces a potentially longer  
25 response latency. In the single-master ring, each slave device may have as many as eight transactions outstanding at any time, as described above.

The remaining combinations share many of the above basic attributes. In a dual-master pair, each master device may initiate read and write operations addressed to the other, and each may have up to eight such transactions  
30 outstanding. No forwarding is required for either device because packets are sent directly to the recipient. A multiple-master ring may contain multiple master

devices and a single slave device. In this configuration, the slave device need not forward packets as all input packets are designated for the single slave device. A multiple-master ring may contain multiple master devices and as many as four slave devices. Each slave device may have up to eight transactions outstanding, and each master device may use some of those transactions. In a preferred embodiment, a master also has the capability to detect a time-out condition or when a response to a request packet is not received. Further aspects of inter-processor communications and configurations are discussed below in connection with FIG. 18.

### Serial Bus

In one preferred embodiment of the invention, the general purpose media processor 12 includes a serial bus (not shown). The serial bus is designed to provide bootstrap resources, configuration, and diagnostic support to the general purpose media processor 12. The serial bus preferably employs two signals, both at TTL levels, for direct communication among many devices. In the preferred embodiment, the first signal is a continuously running clock, and the second signal is an open-collector bi-directional data signal. Four additional signals provide geographic addresses for each device coupled to the serial bus. A gateway protocol, and optional configurable addressing, each provide a means to extend the serial bus to other buses and devices. Although the serial bus is designed for implementation in a system having a general purpose media processor 12, as those skilled in the art will appreciate, the serial bus is applicable to other systems as well.

Because the serial bus is preferably used for the initial bootstrap program load of the general purpose media processor 12, the bootstrap ROM is coupled to the serial bus. As a result, the serial bus needs to be operational for the first instruction fetch. The serial bus protocol is therefore devised so that no transactions are required for initial bus configuration or bus address assignment.

According to the preferred embodiment, the clock signal comprises a continuously running clock signal at a minimum of 20 megahertz. The amount

of skew, if any, in the clock signal between any two serial bus devices should be limited to be less than the skew on the data signal. Preferably, the serial data signal is a non-inverted open collector bi-directional data signal. TTL levels are preferred for communication on the serial bus, and several termination networks may be employed for the serial data signal. A simple preferred termination network employs a resistive pull-up of 220 ohms to 3.3 volts above  $V_{ss}$ . An alternate embodiment employs a more complex termination network such as a termination network including diodes or the "Forced Perfect Termination" network proposed for the SCSI-2 standard, which may be advantageous for larger configurations.

The geographic addressing employed in the serial bus is provided to insure that each device is addressable with a number that is unique among all devices on the bus and which also preferably reflects the physical location of the device. Thus, the address of each device remains the same each time the system is operated. In one preferred embodiment, the geographic address is composed of four bits, thus allowing for up to 16 devices. In order to extend the geographic addressing to more than 16 devices, additional signals may be employed such as a buffered copy of the clock signal or an inverted copy of the clock signal (or both).

The serial bus preferably incorporates both a bit level and packet protocol. The bit level protocol allows any device to transmit one bit of information on the bus, which is received by all devices on the bus at the same time. Each transmitted bit begins at the rising edge of the clock signal and ends at the next rising edge. The transmitted bit value is sampled at the next rising edge of the clock signal. According to one preferred embodiment where the serial data signal is an open collector signal, the transmission of a zero bit value on the bus is achieved by driving the serial data signal to a logical low value. In this embodiment, the transmission of a one bit value is achieved by releasing the serial data signal to obtain a logical high value. If more than one device attempts to transmit a value on the same clock, the resulting value is a zero if any device transmits a zero value, and one if all devices transmit a one value. This provides a "wired-AND" collision mechanism, as those skilled in the art will appreciate. If

two or more devices transmit the same value on the same clock cycle, however, no device can detect the occurrence of a collision. In such cases, the transaction, which may occur frequently in some implementations, preferably proceeds as described below.

5                   The packet protocol employed with the serial bus uses the bit level protocol to transmit information in units of eight bits or multiples of eight bits. Each packet transmission preferably begins with a start bit in which the serial data signal has a zero (driven) value. After transmitting the eight data bits, a parity bit is transmitted. The transmission continues with additional data. A single one  
10 (released) bit is transmitted immediately following the least significant bit of each byte signaling the end of the byte.

                  On the cycle following the transmission of the parity bit, any device may demand a delay of two cycles to process the data received. The two cycle delay is initiated by driving the serial data signal (to a zero value) and releasing  
15 the serial data signal on the next cycle. Before releasing the serial data signal, however, it is preferable to insure that the signal is not being driven by any other device. Further delays are available by repeating this pattern.

                  In order to avoid collisions, a device is not permitted to start a transmission over the serial bus unless there are no currently executing  
20 transactions. To resolve collisions that may occur if two devices begin transmission on the same cycle, each transmitting device should preferably monitor the bus during the transmission of one (released) bits. If any of the bits of the byte are received as zero when transmitting a one, the device has lost arbitration and must cease transmission of any additional bits of the current byte or  
25 transaction.

                  According to the preferred embodiment of the invention, a serial bus transaction consists of the transmission of a series of packets. The transaction begins with a transmission by the transaction initiator, which specifies the target network, device, length, type and payload of the transaction request. The  
30 transaction terminates with a packet having a type field in a specified range. As a result, all devices connected to the serial bus should monitor the serial data signal



to determine when transactions begin and end. A serial bus network may have multiple simultaneous transactions occurring, however, so long as the target and initiator network addresses are all disjoint.

5     Parallel Processing

          In one preferred embodiment of the invention, two or more general purpose media processors 12 can be linked together to achieve a multiple processor system. According to this embodiment, general purpose media processors 12 are linked together using their high bandwidth interface channels 124, either directly or through external switching components (not shown). The dual-master pair configuration described above can thus be extended for use in multiple-master ring configurations. Preferably, internal daemons provide for the generation of memory references to remote processors, accesses to local physical memory space, and the transport of remote references to other remote processors.

10     In a multi-processor environment, all general purpose media processors 12 run off of a common clock frequency, as required by the communication channels 156 that connect between processors.

          Referring to FIG. 18, each general purpose media processor 12 preferably includes at least a pair of inter-processor links 218 (see also FIG. 16(b)). In one configuration, both pairs of inter-processor links 218 can be connected between the two processors 12 to further enhance bandwidth. As shown in FIG. 18(a) several processors 12 may be interconnected in a linear network employing the transponder daemons in each processor. In an alternate embodiment shown in FIG. 18(b), the inter-processor links 218 may be used to join the general purpose media processors 12 in a ring configuration.

20     Alternatively still, general purpose media processors 12 may be interconnected into a two-dimensional network of processors of arbitrary size, as shown in FIG. 18(c). Sixteen processors are connected in FIG. 18(c) by connecting four ring networks. In yet another alternate embodiment, by connecting the inter-processor links 218 to external switching devices (not shown), multi-processors with a large

25     links 218 to external switching devices (not shown), multi-processors with a large

30

42

number of processors can be constructed with an arbitrary interconnection topology.

The requester, responder and transponder daemons preferably handle all inter-processor operations. When one general purpose media processor 12 attempts a load or store to a physical address of a remote processor, the requester daemon autonomously attempts to satisfy the remote memory reference by communicating with the external device. The external device may comprise another processor 12 or a switching device (not shown) that eventually reaches another processor 12. Preferably, two requester daemons are provided each processor 12, which act concurrently on two different byte channels and/or module addresses. The responder daemon accepts writes from a specified channel and module address, which enables an external device to generate transaction requests in local memory or to generate processor events. The responder daemon also generates link level writes to the same external device that communicated responses for the received transaction request. Two such responder daemons are preferably provided; each of which operate concurrently to two different byte channels and/or module addresses.

The transponder daemon accepts writes from a specified channel and module address, which enable an external device to cause a requester daemon to generate a request on another channel and module address. Preferably, two such transponder daemons are provided, each of which act concurrently (back-to-back) between two different byte channel and/or module addresses. As those skilled in the art will appreciate, the requester, responder and transponder daemons must act cooperatively to avoid deadlock that may arise due to an imbalance of requests in the system. Deadlocks prevent responses from being routed to their destinations, which may defeat the benefits of a multi-processor distributed system.

According to one presently preferred embodiment of the invention, the general purpose media processor 12 can be implemented as one or more integrated circuit chips. Referring to FIG. 19, the presently preferred embodiment of the general purpose media processor 12 consists of a four-chip set. In the four-chip set, a general purpose media processor 12 is manufactured as a stand alone

- 42 -  
43

integrated circuit. The stand alone integrated circuit includes a memory management unit 122, instruction and data cache/buffers 118, 120, and an execution unit 100. A plurality of signal input/output pads 260 are provided around the circumference of the integrated circuit to communicate signals to and from the general purpose media processor 12 in a manner generally known in the art.

The second and third chips of the four-chip set comprise in an external memory element 158 and a channel interface device 216. The external memory element 158 includes an interface to the communication channel 156, a cache 150 and a memory interface 152. The channel interface device 216 also includes an interface to the communication channel 156, as well as buffer memory 262, and input/output interfaces 264. Both the external memory element 158 and the channel interface device 216 include a plurality of input/output signal pads 260 to communicate signals to and from these devices in a generally known manner.

The fourth integrated circuit chip comprises a switch 226, which allows for installation of the general purpose media processor 12 in the heterogeneous network 38. In addition to the plurality of input/output pads 260, the switch 226 includes an interface to the communication channel 156. The switch 226 also preferably includes a buffer 262, a router 266, and a switch interface 268.

As those skilled in the art will appreciate, many implementations for the general purpose media processor 12 are possible in addition to the four-chip implementation described above. Rather than an integrated approach, the general purpose media processor can be implemented in a discrete manner. Alternatively, the general purpose media processor 12 can be implemented in a single integrated circuit, or in an implementation with fewer than four integrated circuit chips. Other combinations and permutations of these implementations are contemplated.

There has been described a system for processing streams of media data at substantially peak rates to allow for real time communication over a large heterogeneous network. The system includes a media processor at its core that is capable of processing such media data streams. The heterogeneous network

consists of, for example, the fiber optic/coaxial cable/twisted wire network in place throughout the U.S. To provide for such communication of media data, a media processor according to the invention is disposed at various locations throughout the heterogeneous network. The media processor would thus function  
5 both in a server capacity and at an end user site within the network. Examples of such end user sites include televisions, set-top converter boxes, facsimile machines, wireless and cellular telephones, as well as large and small business and industrial applications.

To achieve such high rates of data throughput, the media processor  
10 includes an execution unit, high bandwidth interface, memory management unit, and pipelined instruction and data paths. The high bandwidth interface includes a mechanism for transmitting media data streams to and from the media processor at rates at or above the gigahertz frequency range. The media data stream can consist of transmission, presentation and storage type data transmitted alone or in a  
15 unified manner. Examples of such data types include audio, video, radio, network and digital communications. According to the invention, the media processor is dynamically partitionable to process any combination or permutation of these data types in any size.

A programmable, general purpose media processor system presents  
20 significant advantages over current multimedia communications. Rather than rigid, costly and inefficient specialized processors, the media processor provides a general purpose instruction set to ease programmability in a single device that is capable of performing all of the operations of the specialized processor combination. Providing a uniform instruction set for all media related operations  
25 eliminates the need for a programmer to learn several different instruction sets, each for a different specialized processor. The complexity of programming the specialized processors to work together and communicate with one another is also greatly reduced. The unified instruction set is also more efficient. Highly specialized general calculation instructions that are tailored to general or special  
30 types of calculations rather than enhancing communication are eliminated.

Moreover, the media processor system can be easily reprogrammed simply by transmitting or downloading new software over the network. In the specialized processor approach, new programming usually requires the delivery and installation of new hardware. Reprogramming the media processor can be  
5 done electronically, which of course is quicker and less costly than the replacement of hardware.

It is to be understood that a wide range of changes and modifications to the embodiments described above will be apparent to those skilled in the art and are contemplated. It is therefore intended that the foregoing detailed  
10 description be regarded as illustrative rather than limiting, and that it be understood that it is the following claims, including all equivalents, that are intended to define the spirit and scope of this invention.